VoIP

DDOS takes on a whole new meaning.

Voice over IP (VoIP) promises to up-end a century-old model of voice telephony by breaking the traditional monolithic service model of the public switched telephone network (PSTN) and changing the point of control and provision from the central office switch to the end user's device. Placing intelligence at the edge, in the Internet tradition, has a number of consequences: a wider community of developers—in particular the large community of Web service developers—can work on voice applications; open interfaces and decomposable functionality facilitate multi-vendor and “homegrown” solutions; and open source and nonproprietary software development can facilitate innovation and experimentation. Users themselves will take a much bigger role in defining, implementing, and controlling the features of telephone services. Contrast this with the more-than-100-year institution of traditional voice service through a single ubiquitous dedicated network, the PSTN. Characteristically, in the PSTN, services are developed hand-in-hand with vendors of the hardware (the circuit switch), the product and its operation are a closely guarded proprietary concern, and, consequently, relatively few vendors of

Security:

Not an Afterthought

VoIP Security:

Not an Afterthought

such products exist and relatively few software engineers are trained in development of voice services.

With its promise of inclusion, innovation, and growth, VoIP also brings challenges. VoIP is not easy to secure. It suffers all of the problems associated with any Internet application, and VoIP security is complicated by its interconnection to the PSTN. A host of trust, implementation, and operational complexities make securing VoIP particularly complex. In fact, the same aspects that make the VoIP software model so powerful—its flexible, open, distributed design—are what make it potentially problematic. There is no central entity, as is the case for the PSTN, responsible for the design, implementation, and monitoring of the voice service. We are moving to an environment where many programmers can create voice applications. But, poorly designed code opens the door to potential security vulnerabilities. Choices regarding the invocation of security mechanisms will be in the hands of the developer, and the testing and validation of these services may be done in a distributed and possibly ad hoc manner. In addition, individual users' expectations and sophistication can and will drive the security of the service much more than in the PSTN. Just as VoIP blurs the line regarding who can develop the service, it blurs the line regarding who can deploy and configure it. Implementers and end users will have an increasing role in configuring and altering voice services, and resulting incidents of misconfiguration will, in many cases, open up security vulnerabilities.

Nonetheless, while VoIP presents a number of security challenges, it also provides an important, fundamental opportunity in security. Placing voice services on a common footing and sharing common protocols with an array of communication and other services also allows us to leverage security solutions across the whole panoply of services. Specifically, much that has been learned (sometimes painfully) about securing the exciting but chaotic world of the Internet applies immediately to VoIP. The same (or similar) classic measures can be applied to VoIP for purposes of confidentiality, integrity, and availability. Economies of scope and scale so achieved should allow us to provide affordable, appropriate, and flexible security

across a range of uses unimaginable in the PSTN world.

In this article, we will consider the security of VoIP in the context of software and network engineering. While we consider the traditional aspects of security, that is, confidentiality, integrity, authentication, authorization, and availability, we will focus on broader issues such as software design, implementation, and interoperability. It is not our intention to provide a recipe book for developing secure VoIP software or a tutorial on implementing security protocols.^{1, 2, 3} Rather, we will consider the directions of VoIP software development and what security challenges and opportunities this might create.

DECOMPOSING VOICE

The key to understanding the effects of VoIP on security is to first understand the decomposition of voice telephony that VoIP engenders. From this, we can examine both the details of its security effects and some of the actions that the software and network engineering com-

Historical Voice Model

FIG 1

The monolithic nature of the pstn voice model.

munities should take.

During the last half of the 20th century, the design of the PSTN progressed from relatively noncomplex mechanical switches, which were hierarchically connected with most of the intelligence and control residing locally at the switch, to complex digital switches handling many functions, still with most of the intelligence and control local, to the present “intelligent network” design, wherein a significant amount of the intelligence and control has migrated away from the local switch into the core of the network. This transition included a move toward digital systems, software control, and general-purpose computing platforms. As shown in figure 1, resource control migrated from the local switch to deeper within the network, but always within the control of the carrier. In moving intelligence toward the core of the network, the PSTN service providers were looking (among other things) to improve efficiency and enhance security.

As described elsewhere in this issue (see, Sherburne and Fitzgerald’s “You Don’t Know Jack about VoIP” on page 30 and James E. Coffman’s “Not Your Father’s PBX” on page 40), there are a wide variety of possible VoIP implementations, each depending on the selection of protocols and the specific implementation of those protocols. Services can now be offered in a highly distributed manner, with the various functions cobbled together by the end user or by a service aggregator—and maintained

throughout the network and controlled in various ways by users, network providers, or both. This diversity of providers, components, and configurations (as depicted in figure 2) will be among the implementation challenges facing the security of VoIP. Furthermore, the character of VoIP offerings in the market is much more diverse than traditional telephony offerings have been, driven in part by experimentation with business models at a similar pace to experimentation with technical possibilities. As in wireless telephony, these VoIP models are demonstrating that users are willing to sacrifice traditional expectations, such as quality-of-service or availability, for some other attribute such as price point or mobility. This does not mean that users don’t value quality-of-service or availability, but it does undermine the concept matured in the PSTN that there is a sharp transition between unacceptable and acceptable service, regardless of other tradeoffs. This ability and willingness to differentiate—in terms of service, protocols, implementation, provider, and price—has significant implications for the security models that will develop.

VOIP AND SECURITY

Just as VoIP can be deployed in fairly safe environments (e.g., secure and/or trusted), it can also be deployed in potentially unsafe environments. In such an environment, the security vulnerabilities and potential attacks are

numerous. Attacks can be launched on the underlying network, the transport protocols, the VoIP devices (e.g., servers and gateways), the VoIP application, other related applications (e.g., Dynamic Host Configuration Protocol, DHCP), the underlying operating systems, and more. Vulnerabilities introduced in the design of the product will be a continuous concern, as will improper configuration and implementation. In fact, the fundamental paradigm shift—voice as just another service running on a shared, IP-based infrastructure—calls into question our basic expectations of the security of

VoIP Security:

Not an Afterthought

voice. The features of the technology employed often subtly, or not so subtly, color security expectations and concerns. For example, the classic attack on confidentiality prior to the electronic age was “steaming” an envelope open; risks were to be expected and countermeasures were developed with this particularly in mind. But this attack and its countermeasures are clearly irrelevant to a telephone conversation. What survives is a common concern about confidentiality, and it is to such basic security concerns and system characteristics that we turn in Table 1 and the text that follows.

SECURITY CONCERNS

Expectations of security in the wired PSTN are strongly influenced by its physical character and operation. Privacy and confidentiality are aided by the difficulty in physically accessing wires in order to tap them. While it may be easy for anyone to tap a telephone line at the side of a house (if they can gain covert physical access), it is much harder to do so within the network. Consequently, users feel relatively secure against all but determined attackers and rarely consider methods to further protect confidentiality. Law enforcement, relying on this expecta-

tion and backed by the Communications Assistance for Law Enforcement Act (CALEA), can expect to intercept and listen to conversations when authorized by a court order. But with VoIP, the voice conversation can flow over many different physical networks with different interception characteristics. The user, less confident now in the consistent security of the network, may invoke end-to-end security mechanisms such as strong encryption. Law enforcement agents, in turn, now have a much greater challenge when they want to intercept traffic, both from the diverse routing possible and from strong encryption. This new interplay between privacy and other social policy goals is currently creating substantial controversy. To examine this issue, we might distinguish between different flavors of VoIP: carrier and non-carrier VoIP. With non-carrier-flavored VoIP, it appears that CALEA is presently not applicable—in essence, such VoIP is like e-mail, and CALEA expressly disclaims any application to information services. With more carrier-like flavors, and particularly those with a connection to the PSTN, the issue is now posed in a recent FBI (Federal Bureau of Investigation) petition to the Federal Communications Commission (FCC). The FCC may grant this petition, concluding

TABLE 1 Security Concerns and System Characteristics

	Wired PSTN Measures	VoIP Measures
Security Concerns		
Confidentiality	Physical security	Encryption techniques
Integrity	Physical security	Encryption techniques
Availability	Physical access control	Network/Service access control
Authentication	Physical connectivity, voice recognition, caller ID	Login, password
Authorization	Caller ID, access control	Access control, role-based authorization
User Expectation	Assumed and static	Variable
Implementation and Design Concerns		
Software design	Large, monolithic, complex	Variable, distributed, complex
Interoperability	Centralized and tested	Distributed and potentially ad hoc
Software implementation	Centralized and tested	Distributed and potentially ad hoc

that VoIP is a telecom service for the purposes of CALEA. The wireless PSTN—mobile telephony—has already seen some changes similar to what VoIP engenders. For example, the ease of eavesdropping on a radio transmission relative to physically connecting to wires pushed the cellular industry to develop and standardize encryption techniques to protect the privacy of cellular phone conversations. But VoIP will go substantially farther in changing how security is conceived and addressed.

Integrity is the idea that we can rely on what we send or receive to be as we intended. With VoIP traveling over a variety of networks, it is conceivable that messages (voice or signaling) could be modified or spoofed within the network, although the human capability of speaker recognition makes this more challenging than, say, spoofing an e-mail. The natural response is, again, a shift toward security using encryption and related techniques. The distribution of keys to enable these types of security mechanisms will likely draw heavily on continued public key infrastructure developments. In particular, current mechanisms for distributing certificates and establishing trust may work well when we want to authenticate a small number of important central organizations, but use of the public key infrastructure to authenticate all VoIP clients creates problems. For example, how exactly can we be sure that each individual who applies for and pays for a certificate is the person who they claim to be?

Availability means that the service is there when desired. In the PSTN, availability was ensured through network engineering and limited access to the network. With voice traveling over an Internet-based network, issues such as denial of service (DoS) and distributed DoS attacks represent a significant threat to the availability of the services. Addressing these issues will require adoption of the same techniques being applied elsewhere in the Internet, namely appropriate network design and access control. Of course, DoS attacks occurred in the PSTN as well through such mechanisms as “wardialers.” A disincentive to launching DoS attacks is the strong authentication technology built into the PSTN. The inability to identify the instigator of a DoS attack on the Internet makes these attacks more difficult to counter. While research is underway to increase the ability to trace such attacks, these efforts have met with only limited success. However, it is important to stress that the extreme distribution of VoIP and the lack of monolithic bottlenecks speak to its greater long-term strength in meeting both DoS attacks and congestion threats.

Classic security concerns also include being sure of who is involved in a communication (authentication)

and only providing them with appropriate capabilities (authorization). The phone company identifies customers based on the line over which they are connecting. The PSTN’s physical characteristics—wire based, circuit switching based on a telephone number as the address—mean that authentication of a receiver is limited to reliance on the circuit switch correctly routing the call to the terminal identified by the telephone number, while authentication and authorization of a sender is restricted to decisions that can be made based on the telephone number of the sender (through caller ID) and recognizing the caller’s voice. Unfortunately, voice recognition proves ineffectual on many important calls (e.g., 911 emergency calls and business transactions) involving parties who have never spoken. Here, VoIP offers the potential for improved capabilities, as authentication techniques can be applied (e.g., user id and password, biometrics, and physical devices such as smart cards and keys) so that we authenticate not a device but rather a person. Moreover, much richer authorization capabilities also exist to route a call differently based on authenticated attributes of the sender or receiver (e.g., my family members’ calls follow me wherever I go, but business associates can only reach

**Expectations of security
in the wired PSTN are
strongly influenced
by its physical character
and operation.**

me at the office or on my cell phone during business hours). But the combinatorics of possible actions based on user identity could easily become unmanageable. Here, the key developments are in role-based authorization, in which capabilities are determined based on the user’s membership in a class and trust may be established between entities sufficiently to authorize users across entity boundaries based on commonly understood roles.

Implementation of these more advanced authentication and authorization techniques should help address such looming issues as VoIP spam and the host of general Internet ills that can be expected to afflict VoIP applications wherever their creators find an opportunity (e.g., spyware applications that record a user’s behavior—in this case perhaps the people they call—and harvest the

VoIP Security:

Not an Afterthought

data unbeknownst to the user). Again, general improvements in Internet security will help secure VoIP as the complement to the way that integration into the Internet exposes VoIP to these problems. On the legal side, we are only beginning to see hints of what might happen. It is unclear whether VoIP spam would be subject to the same regulations as other e-mail spam, which, at present, are minimal, or if it will follow the more aggressive regulation seen in the traditional telephony space.

Finally, one consequence of conjoining domains with very different security assumptions—the open and semi-chaotic Internet and the closed and isolated PSTN—is the creation of opportunities for leakage of threats from robust into vulnerable networks. An unsecured gateway used for translating signaling between VoIP and PSTN signaling could introduce threats to elements of the PSTN that were not designed with such threats in mind. The solution here is a higher level of care in gateway functions and efforts to provide a higher degree of isolation than would be a normal design concern in internal nodes of either network.

DESIGN AND IMPLEMENTATION CONCERNS

The way VoIP systems are developed, the range of functionality they can encompass, and the systems in which they are embedded have important implications for security.

For VoIP software development, key concerns include software stability, robustness, and interoperability. These issues influence the ultimate security of the software, in that flaws, instability, lack of robustness, and lack of interoperability all create potential security breaches—a fact well understood by the creators of operating systems and applications on the Internet. In addressing these issues, developers will need to consider the same principles that exist in the data world as they develop robust and secure VoIP applications. The software community has exerted considerable effort in developing (if not always following) sound software engineering practices. Being built on the same protocols and infrastructure as other Internet applications, VoIP will be subject to all of the security issues that we face on the Internet, including

viruses, denial of service, software exploits, spam, and unauthorized access. The good news is that, while never a solved problem, the process of securing the Internet is fairly mature. While VoIP developers should heed the security lessons of the Internet, there are lessons from the PSTN to consider as well. In terms of operation, there are classic telephony concerns such as feature interaction. Feature interaction occurs when multiple features interact in a way that impedes the desired operation of the service, such as when a user forwards a call to a number that has blocked that forwarded number. However, as described by Lennox and Schulzrinne,⁴ while VoIP presents a new set of complications for feature interaction, it also provides a number of new solutions.

Most VoIP implementations separate voice transport (e.g., Realtime Transport Protocol, RTP), signaling (e.g., the Session Initiation Protocol, SIP), and service creation from one another, and allow each to maintain a fairly complex set of options. These options are not static but are intentionally allowed to evolve. To further complicate matters, many protocols (both proprietary and non-proprietary) exist in the VoIP space, some of which are complementary while others are somewhat substitutable (for example, there are two widely known VoIP signaling protocols—the SIP and H.323 standards). The inherent complexity this creates mirrors that of the Internet and demands the same types of interoperability mechanisms that have evolved there (e.g., formal standards bodies; industry consortia; one-on-one interoperability testing; sharing of common, well-debugged module implementations; and carefully structured fallback mechanisms—the ability to fall back to a less desirable but functional behavior that can be implemented by all parties or the ability to download functionality when required). For instance, the SIP community participates in interoperability tests through SIPit events (see <http://www.sipit.net>). These events allow implementers to test the interoperability of their products, including various security functions. To further aid implementation, the SIP specification provides a separate implementers' section, which includes a description of the various security methods, requirements, and protocols.

To briefly explore how security is provided in one such VoIP-related protocol, let's consider SIP. In the Internet tradition, SIP relies heavily on other protocols to provide other functions, including security. SIP uses IPSec (Internet Protocol Security, a network layer security protocol) and TLS (Transport Layer Security) a transport layer security protocol) and borrows from existing methods such as S/MIME (Secure Multipurpose Internet Email Extensions) and HTTP Digest Authentication. SIP does not provide the encryption for the media (Real Time Protocol, RTP) stream; this is addressed separately as described in the RTP specification³ and through such techniques as Secure RTP.² A non-exhaustive list of SIP security protocols is outlined in table 2.

Analogous to the shift in software concerns from those of monolithic code development to distributed development is the shift in network robustness from a centralized, hierarchical system to a distributed system. Both have their own vulnerabilities. Characteristic of centralized systems is vulnerability to component failure, even with robustness as a design feature—a characteristic sharply illuminated when destruction of a key central office in Manhattan disabled local wireline phone service during the 9/11 attacks. But distributed systems have their own robustness vulnerabilities, an example being the difficult-to-defeat distributed denial of service attack used by hackers to overload servers. VoIP moves voice network robustness and control concerns sharply from the former to the latter.

The shift in user expectations toward much greater flexibility and control creates issues that go beyond software development. Indeed, many VoIP problems, while not necessarily separate from the software, are architectural or configuration problems. Where the architecture or configuration impedes the operation of VoIP, such as the problem of dynamic port allocation and the use of firewalls, both software and network engineers need to consider methods of addressing the problem

while recognizing that security may complicate the environment. Another significant challenge to overcome is the testing of VoIP services in a highly distributed and diverse environment. While a number of tools exist for testing in homogeneous VoIP environments (e.g., sipsak is a simple open source tool for testing SIP applications; a number of open source test tools, and other VoIP applications, can be found at the VoIP-Info Web site: <http://www.voip-info.org/wiki-Open+Source+VOIP+Software>), few tools exist to examine end-to-end operation in environments with mixes of proprietary and non-proprietary technologies, numerous protocols, and interconnection with the PSTN.

In this environment, users will have the potential for a much more active role in their voice service. Nonetheless, while a user could misconfigure their service, they now also have the option of fixing it. This ability to intervene should make an interesting case study in user adoption of technology and the evolution of cultural expectations of technology. (As we have commented, mobile telephony shows these are not as static as many PSTN veterans assume, and the Internet has proven to be fertile ground for developing brand-new user behaviors). Some users will seek to intervene and, as such, software engineers should think about how to enable users by providing them with tools to test and manage their own voice service.

CONCLUSION

VoIP presents a number of interesting security challenges that differ substantially from those of traditional telephony. In addressing these challenges, we might consider the roles of the vendor, service provider, and implementer communities.

Vendor. Software products will more frequently consist of multiple-vendor modular solutions written in standard languages (as opposed to more monolithic vendor internal programs). To complicate issues, the software will be in a constant state of evolution (mirroring the rest

of the computer software industry). Further, software (following VoIP protocol bloat) will continue to grow in size and complexity, and vendors will have to wrestle with issues of differentiation through new integrated services, proprietary extensions, and closed code. Vendors will need to adapt many

TABLE 2 Examples of SIP (Session Initiation Protocol)-Related Security Protocols

Protocol	Source
Transport Layer Security (TLS)	RFC 2246 and updates in RFC 3546
Internet Security (IPSec)	RFC 2401
Digest authentication	Parallels RFC 2617 (as described in RFC 3261)
S/MIME	RFCs 1847, 2630, and 2633 (as described in RFC 3261)
Privacy mechanisms	RFC 3323

VoIP Security:

Not an Afterthought

practices that are common in computing and Internet products to VoIP applications, including (1) intentional investment in interoperability efforts with other vendors at multiple levels, (2) exploiting technical opportunities available in an IP-based network but in “fallback-friendly” mode that protects some functionality even when advanced features will not work, (3) continuing and perhaps increasing effort devoted to software engineering practices intended to support predictable and secure performance in realtime, critical systems, (4) leveraging a broad base of Web-savvy programmers for both ideas and programming effort, and (5) participating in the debate about social policy and regulatory reform around VoIP.

Provider. Service providers face tremendous diversity in potential offerings and configuration of their infrastructure, creating both security opportunities and risks. To address this, they will need to ensure their internal security expertise and develop the ability to securely integrate modular components from multiple vendors through constant testing and evaluation. Security will be a cooperative venture involving the service provider, vendors, and customers. Finally, service providers should consider participating in the debate about social, policy, and regulatory reform, as these issues will have a profound impact on what defines the security of VoIP.

User. These individuals will face a larger (possibly bewildering) array of different service offerings. Sophisticated users (i.e., users for whom telecommunications is important enough to warrant a substantial investment of their time) will want to increase their awareness about the best tradeoffs for them among cost, quality, and security; less sophisticated users will rely on trusted brands or on government and industrial certification. Finally, we are entering a period in which users can examine, develop, and alter their voice service, which is a radical departure from the past. Ultimately, an end user interested in maintaining a high level of security has the ability to participate in creating this environment. Q

REFERENCES

1. Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and Schooler, E. SIP: Session Initiation Protocol, Internet RFC 3261.
2. Baugher, M., McGrew, D., Naslund, M., Carrara, E., and Norrman, K. The Secure Real-time Transport Protocol (SRTP), Internet RFC 3711.
3. Schulzrinne, H., Casner, S., Frederick, R., and Jacobson, V. RTP: A Transport Protocol for Real-Time Applications, Internet RFC 3550.
4. Lennox, J., and Schulzrinne, H. Feature Interaction in Internet Telephony. Sixth Workshop on Feature Interactions in Telecom and Software Systems, Glasgow, Scotland, June 2000.

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

DOUGLAS C. SICKER (douglas.sicker@colorado.edu) is an assistant professor in the computer science department at the University of Colorado at Boulder, with a joint appointment in the Interdisciplinary Telecommunications Program. Prior to this he was chief of the Network Technology Division in the Office of Engineering and Technology at the Federal Communications Commission. His research interests include network applications and security, and public policy. Sicker holds a B.S., M.S., and Ph.D. from the University of Pittsburgh, and is a senior member of the IEEE and a member of ACM and the Internet Society.

TOM LOOKABAUGH (tom.lookabaugh@colorado.edu) is an assistant professor in the computer science department at the University of Colorado at Boulder, as well as faculty director of Interdisciplinary Telecommunications. Prior to joining the university, he spent 13 years in high-technology businesses in Silicon Valley, including cofounding an entertainment video compression manufacturer, DiviCom. His research interests include ubiquitous networking, multimedia, and computer security. He holds a B.S. in engineering physics from Colorado School of Mines, and M.S. degrees in electrical engineering, engineering management, and statistics and a Ph.D. in electrical engineering from Stanford University. Lookabaugh is a member of ACM and the Academy of Management.

© 2004 ACM 1542-7730/04/0900 \$5.00